

Stochastic Transition Systems^{*}

Luca de Alfaro

University of California at Berkeley
`dealfaro@eecs.berkeley.edu`

Abstract. Traditional methods for the analysis of system performance and reliability generally assume a precise knowledge of the system and its workload. Here, we present methods that are suited for the analysis of systems that contain partly unknown or unspecified components, such as systems in their early design stages.

We introduce stochastic transition systems, a high-level formalism for the modeling of timed probabilistic systems. Stochastic transition systems extend current modeling capabilities by enabling the representation of transitions having unknown delay distributions, alongside transitions with zero or exponentially-distributed delay. We show how these various types of transitions can be uniformly represented in terms of nondeterminism, probability, fairness and time, yielding efficient algorithms for system analysis. Finally, we present methods for the specification and verification of long-run average properties of STSs. These properties include many relevant performance and reliability indices, such as system throughput, average response time, and mean time between failures.

1 Introduction

The analysis of system performance and reliability is an essential part of the design of many computing and communication systems. Most approaches to the computation of performance and reliability indices presuppose that the structure of the system is known in detail, and that the values of the transition probabilities and the delay distributions are precisely known. Here, we describe methods that are suited to the evaluation of systems that are still in their early stages of design, when not all the system components may have been designed, and when relevant quantities may be known only with some approximation.

We introduce *stochastic transition systems* (STSs), a high-level modeling language for timed probabilistic systems. Stochastic transition systems provide a concise and compositional way to describe the behavior of systems in terms of probability, waiting-time distributions, nondeterminism, and fairness. In particular, the execution model of STSs extends that of generalized stochastic Petri

^{*} This work was partially supported by the NSF grant CCR-95-27927, by DARPA under the NASA grant NAG2-892, by the ARO grant DAAH04-95-1-0317, by ARO under the MURI grant DAAH04-96-1-0341, by Army contract DABT63-96-C-0096 (DARPA), by the ONR YIP award N00014-95-1-0520, by the NSF CAREER award CCR-9501708, and by the NSF grant CCR-9504469.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 1998		2. REPORT TYPE		3. DATES COVERED 00-00-1998 to 00-00-1998	
4. TITLE AND SUBTITLE Stochastic Transition Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California at Berkeley, Department of Electrical Engineering and Computer Sciences, Berkeley, CA, 94720-1770				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

nets [ABC84] and of stochastic process algebras such as TIPP [GHR93], PEPA [Hil96] and EMPA [BG96] with the introduction of nondeterminism and of transitions with unspecified delay distribution. These features enable the modeling of unknown (or imprecisely known) arrival rates and transition probabilities, as well as the modeling of schedulers with unspecified behavior.

We provide two semantics for STSs. The first one is an informal semantics that can be used to gain an intuitive understanding of STSs, and to guide the construction of system models. The second semantics is defined by providing a translation from STSs to *fair timed probabilistic systems* (fair TPSs), a low-level computational model based on Markov decision processes that is well suited to the application of verification algorithms. The relation between an STS and its translation TPS parallels the relation between a first-order transition system and its representation as a state-transition graph; in particular, the state space of the translation TPS coincides with that of the STS. We show that the translation precisely captures the informal semantics of STSs, justifying the use of the informal semantics in the construction of system models.

The translation relies on a new notion of fairness for probabilistic systems, called *probabilistic fairness*. Unlike previous notions of fairness, which refer to the transitions that are enabled and taken along system behaviors [Var85, MP91, KB96], probabilistic fairness is a structural condition on the policies that govern the resolution of nondeterministic choices. The condition states that, for every policy, there must be a *fixed* $\varepsilon > 0$ such that every fair alternative is selected with probability at least ε . Probabilistic fairness enables the faithful representation of transitions with unspecified delay distributions. Probabilistic fairness also simplifies the analysis of several algorithms, since its basic ingredients —policies and probability— are already present in Markov decision processes.

We then turn our attention to the specification and verification of *long-run average properties* of probabilistic systems. Long-run average properties refer to the average behavior of a system, measured over a period of time whose length diverges to infinity. In a purely probabilistic system, these properties are related to the steady-state distribution of the Markov chain corresponding to the system. We specify long-run average properties of systems by attaching labels to the system states and transitions, following a simplified version of the approach of [dA98]. The labels specify system tasks, whose long-run average outcome or duration can be measured. This enables the specification of several reliability and performance indices, such as throughput, average response time, and mean time between failures.

Finally, we present algorithms for verifying that the performance and reliability specifications of an STS are met even under the most unfavorable combination of nondeterministic behavior and choice of delays for the transitions with unknown delay distributions. The verification process is based on an adaptation of the algorithms presented in [dA98] to systems that include fairness. We show that the presence of fairness does not increase the complexity of the verification problem, which can again be solved in polynomial time in the size of the fair TPS. The analysis of the verification algorithms also shows that, when consider-

ing long-run average properties of finite-state systems, our notion of probabilistic fairness yields the same verification algorithms as the *weak fairness* of [KB96], showing that the two notions are equivalent in this context.

2 Stochastic Transition Systems

Stochastic transition systems (STSs) have been inspired by the *fair transition systems* of [MP91] and by the *real-time probabilistic processes* of [ACD92]. A *stochastic transition system* (STS) is a triple $\mathcal{S} = (\mathcal{V}, \Theta, \mathcal{T})$, where:

- \mathcal{V} is a finite set of typed *state variables*, each with finite domain. The (finite) state space S consists of all type-consistent interpretations of the variables in \mathcal{V} . We denote by $s[x]$ the value at state $s \in S$ of $x \in \mathcal{V}$; the interpretation function $\llbracket \cdot \rrbracket$ is extended to terms in the obvious way.
- Θ is an assertion over \mathcal{V} denoting the set $\{s \in S \mid s \models \Theta\}$ of *initial states*.
- \mathcal{T} is a set of *transitions*.

With each transition $\tau \in \mathcal{T}$ are associated the following quantities:

- An assertion \mathcal{E}_τ over \mathcal{V} , which specifies the set of states $\{s \in S \mid s \models \mathcal{E}_\tau\}$ on which τ is enabled.
- A number $m_\tau > 0$ of *transition modes*. Each transition mode $i \in \{1, \dots, m_\tau\}$ corresponds to a possible outcome of τ , and is specified by:
 - A set of assignments $\{x' := f_{i,x}^\tau\}_{x \in \mathcal{V}}$, where each $f_{i,x}^\tau$ is a term over \mathcal{V} . These assignments define the function $f_i^\tau : S \mapsto S$, which maps every state $s \in S$ into a successor $s' = f_i^\tau(s)$ such that $s'[x] = s[f_{i,x}^\tau]$ for all $x \in \mathcal{V}$.
 - The probability $p_i^\tau \in [0, 1]$ with which mode i is chosen. We require $\sum_{i=1}^{m_\tau} p_i^\tau = 1$.

The set \mathcal{T} of transitions is partitioned into the two subsets \mathcal{T}_i and \mathcal{T}_d of *immediate* and *delayed* transitions. Immediate transitions must be taken as soon as they are enabled. A subset $\mathcal{T}_f \subseteq \mathcal{T}_i$ indicates the set of *fair* transitions. In turn, the set \mathcal{T}_d of delayed transitions is partitioned into the sets \mathcal{T}_e and \mathcal{T}_u , where:

- \mathcal{T}_e is the set of transitions with *exponential delay distribution*. With each $\tau \in \mathcal{T}_e$ is associated a transition rate $\gamma_\tau > 0$.
- \mathcal{T}_u is the set of transitions with *unspecified delay distributions*. These transitions are taken with non-zero delay, but the probability distribution of the delay, and the possible dependencies between this distribution and the system's present state or past history are not specified.

Given a state $s \in S$, we indicate by $\mathcal{T}(s) = \{\tau \in \mathcal{T} \mid s \models \mathcal{E}_\tau\}$ the set of transitions enabled at s . To insure that $\mathcal{T}(s) \neq \emptyset$ for all $s \in S$, we implicitly add to every STS an *idle transition* $\tau_{idle} \in \mathcal{T}_e$, defined by $\mathcal{E}_{\tau_{idle}} = \text{true}$, $m_{\tau_{idle}} = 1$, $p_1^{\tau_{idle}} = 1$, $\gamma_{\tau_{idle}} = 1$, and by the set of assignments $\{x' := x\}_{x \in \mathcal{V}}$. The choice of an unitary transition rate is arbitrary.

2.1 Informal Semantics of Stochastic Transition Systems

We present here an informal semantics of STSs, which can be used to gain an intuitive, but accurate, understanding of their behavior. In a later section, we show that this informal semantics precisely corresponds to the formal semantics, defined by translation into lower-level computational models.

In the informal semantics, the temporal evolution of the system state is represented by a *timed trace*. A timed trace is an infinite sequence $(s_0, I_0), (s_1, I_1), \dots$ of pairs, where $I_k \subseteq \mathbb{R}^+$ is a closed interval and s_k is a system state, for $k \geq 0$. The intervals must be contiguous, i.e. $\max I_k = \min I_{k+1}$ for all $k \geq 0$, and the first interval must begin at 0, i.e. $\min I_0 = 0$. A pair (s_k, I_k) in a timed trace indicates that during the interval of time I_k the system is in state s_k . The choice of considering only closed intervals is arbitrary. Note that point intervals are permitted: they represent transitory states in which an immediate transition is taken before time advances. These transitory states are very similar to the *vanishing markings* of *generalized stochastic Petri nets* (GSPNs) [ABC84].

The initial state s_0 of a timed trace must satisfy $s_0 \models \Theta$. For $k \geq 0$, state s_k determines the expected duration of I_k and the next state s_{k+1} as follows:

- *Some immediate transition enabled.* If $\mathcal{T}(s_k) \cap \mathcal{T}_i \neq \emptyset$, then the duration of I_k is 0. A transition $\tau \in \mathcal{T}(s_k) \cap \mathcal{T}_i$ is chosen nondeterministically, subject to fairness requirements: if $\tau \in \mathcal{T}_f$, then τ must be chosen with non-zero probability.

Once τ has been chosen, each transition mode $i \in [1..m_\tau]$ is chosen with probability p_i^τ , and the successor state is given by $s_{k+1} = f_i^\tau(s)$.

- *Only delayed transitions enabled.* If $\mathcal{T}(s_k) \subseteq \mathcal{T}_d$, let $\mathcal{T}_e(s_k) = \mathcal{T}(s_k) \cap \mathcal{T}_e$ and $\mathcal{T}_u(s_k) = \mathcal{T}(s_k) \cap \mathcal{T}_u$. The transition rates γ_τ for $\tau \in \mathcal{T}_e(s)$ are given; we select nondeterministically $\gamma_\tau > 0$ for $\tau \in \mathcal{T}_u(s_k)$. The expected duration of I_k is then given by $1 / \sum_{\tau \in \mathcal{T}(s_k)} \gamma_\tau$, and each transition $\tau \in \mathcal{T}(s)$ is chosen with probability $\gamma_\tau / \sum_{\tau' \in \mathcal{T}(s_k)} \gamma_{\tau'}$.

Once τ has been chosen, each transition mode $i \in [1..m_\tau]$ is chosen with probability p_i^τ , and the successor state is again given by $s_{k+1} = f_i^\tau(s)$.

Time divergence. In our definition of timed trace, we have not ruled out the possibility of traces along which time does not diverge. These traces can arise, since the time intervals in the trace can be point intervals, or can be arbitrarily small. In a later section, we provide a method for checking that non-time-divergent traces occur with probability 0.

2.2 An Example of STS

As a simple example of STS, we consider a model for a system consisting of a commuter that continually travels between cities A and B, each way passing through an intermediate city C. Cities A and C are connected by highway link 1, cities C and B by link 2. Each link can be in good conditions, in poor conditions, or undergoing repair: for $i = 1, 2$, the state of link i is represented by variable l_i , with domain $\{g, p, r\}$. For each link, the transition from *good* to *poor* has

rate $\gamma_{gp} = 0.05$, and the transition from *repair* to *good* has rate $\gamma_{rg} = 0.1$. The transition from *poor* to *repair* has unspecified delay distribution: the scheduling of road repairs follows criteria that are not known to the layperson.

The commuter can be at one of 4 states, depending on which segment must be traversed next and in which direction. The state of the commuter is represented by variable c , with domain $\{1, 2, 3, 4\}$: we let $c = 1$ when $A \rightarrow C$ is the next trip to be undertaken, and similarly $c = 2$ for $C \rightarrow B$, $c = 3$ for $B \rightarrow C$, and $c = 4$ for $C \rightarrow A$. Depending on the conditions of the next link, the commuter traverses the link with rate $\gamma_g = 0.5$, $\gamma_p = 0.3$, or $\gamma_r = 0.1$.

The STS $\mathcal{S} = (\mathcal{V}, \Theta, \mathcal{T})$ has variables $\mathcal{V} = \{c, l_1, l_2\}$ and initial condition $\Theta : c = 1 \wedge l_1 = g \wedge l_2 = g$. The set of transitions is $\mathcal{T} = \{\tau_{gp,i}, \tau_{pr,i}, \tau_{rg,i}\}_{i=1,2} \cup \{\tau_g, \tau_p, \tau_r\}$, where transition $\tau_{gp,i}$ models link i going from good to poor, transition τ_g models the commuter traversing a good link, and the meaning of the other transitions can be analogously inferred. We list only a few representative transitions; the others are similar. For brevity, while describing transition τ we write \mathcal{E} instead of \mathcal{E}_τ , and so forth.

- For $i = 1, 2$, $\tau_{gp,i} \in \mathcal{T}_e$ is defined by $\mathcal{E} : l_i = g$; and $\gamma = 0.05$; $m = 1$; $p_1 = 1$; and $l'_i := p$, $l'_{3-i} := l_{3-i}$, $c' := c$.
- For $i = 1, 2$, $\tau_{pr,i} \in \mathcal{T}_u$ is defined by $\mathcal{E} : l_i = p$; and $m = 1$; $p_1 = 1$; and $l'_i := r$, $l'_{3-i} := l_{3-i}$, $c' := c$.
- $\tau_g \in \mathcal{T}_e$ is defined by $\mathcal{E} : [(c = 1 \vee c = 4) \wedge l_1 = g] \vee [(c = 2 \vee c = 3) \wedge l_2 = g]$; and $\gamma = 0.5$; $m = 1$; $p_1 = 1$; and $c' := (c \bmod 4) + 1$, $l'_1 := l_1$, $l'_2 := l_2$.

Alternatively, consider the case in which links in poor conditions are scheduled for repair with rate *at least* 0.1. To model this case, it is possible to introduce additional transitions $\tau'_{pr,i} \in \mathcal{T}_e$ for $i = 1, 2$. These transitions are defined like $\tau_{pr,i}$, $i = 1, 2$, except that they have rate $\gamma = 0.1$. More complex combinations of exponential-delay and unspecified-delay transitions can be used to model more general types of partial knowledge about transition rates.

2.3 Related Models for Probabilistic Systems

Stochastic transition systems are related to several other models for probabilistic systems. The execution model of STSs is related to that of generalized stochastic Petri nets (GSPNs) [ABC84]. In particular, STSs generalize GSPNs by introducing transitions with unspecified delay distributions, and by introducing the possibility of nondeterministic choice among enabled immediate transitions. STSs extend in a similar way also the *probabilistic finite-state programs* of [PZ86] and the *real-time probabilistic processes* of [ACD92]. The introduction of nondeterminism and of transitions with unspecified delay distributions, and the capability to deal with these features in the verification process, also represents an innovation with respects to probabilistic process algebras for performance modeling, such as TIPP [GHR93], PEPA [Hil96] and EMPA [BG96]. *Probabilistic automata* [SL94, Seg95] are another model that has been proposed for probabilistic real-time systems. Probabilistic automata are more closely related to *timed probabilistic systems*, our low-level model of computation, than to STSs.

3 Translating STSs into Low-Level System Models

The formal semantics of STSs is defined by translating STSs into *fair timed probabilistic systems* (fair TPSs), a low-level computational model based on Markov decision processes. Besides providing us with a formal semantics for STSs, the translation is also used in the verification process, since the verification algorithms will be applied to the fair TPSs obtained by translating the STSs.

3.1 Timed Probabilistic Systems

A *Markov decision process* (MDP) is a generalization of a Markov chain in which a set of possible actions is associated with each state. To each state-action pair is associated a probability distribution, used to select the successor state [Der70]. We consider a fixed set of typed state variables \mathcal{V} , coinciding with the variables of the STS. An MDP $\Pi = (S, A, p)$ consists of the following components:

- A finite set S of states, where each $s \in S$ assigns value $s[[x]]$ to each $x \in \mathcal{V}$.
- For each $s \in S$, $A(s)$ is a non-empty finite set of *actions* available at s .
- For each $s, t \in S$ and $a \in A(s)$, $p_{st}(a)$ is the probability of a transition from s to t when action a is selected. For every $s, t \in S$ and $a \in A(s)$, we have $0 \leq p_{st}(a) \leq 1$ and $\sum_{t \in S} p_{st}(a) = 1$.

A *behavior* of an MDP is an infinite sequence $\omega : s_0 a_0 s_1 a_1 \dots$ of alternating states and actions, such that $s_i \in S$, $a_i \in A(s_i)$ and $p_{s_i, s_{i+1}}(a_i) > 0$ for all $i \geq 0$. For $i \geq 0$, the sequence is constructed by iterating a two-phase selection process. First, an action $a_i \in A(s_i)$ is selected nondeterministically; second, the successor state s_{i+1} is chosen according to the probability distribution $p_{s_i, s_{i+1}}(a_i)$. A *timed probabilistic system* (TPS) $\Pi = (S, A, p, S_{in}, time)$ consists of an MDP (S, A, p) , and of the following additional components [dA97a, dA98]:

- A subset $S_{in} \subseteq S$ of initial states. Each behavior of Π must begin with a state in S_{in} .
- A labeling *time* that associates to each $s \in S$ and $a \in A(s)$ the *expected* amount of time $time(s, a) \in \mathbb{R}^+$ spent at s when action a is selected.

We will often associate with an MDP or TPS additional labelings; the labelings will be simply added to the list of components. We define the *size* of an MDP or TPS Π to be the length (in bits) of its encoding, where we assume that transition probabilities are encoded as the ratio between integers.

To be able to assign probabilities to sets of behaviors, we need to specify the criteria used to choose the actions. To this end, we use the concept of *policy* [Der70], closely related to the adversaries of [SL94, Seg95] and to the schedulers of [Var85, PZ86]. A policy η is a set of conditional probabilities $Q_\eta(a \mid s_0 s_1 \dots s_n)$, for all sequences of states $s_0 s_1 \dots s_n \in S^+$ and all $a \in A(s_n)$. The conditional probability $Q_\eta(a \mid s_0 s_1 \dots s_n)$ is the probability with which action $a \in A(s_n)$ is chosen after the system has followed the sequence of states $s_0 s_1 \dots s_n$. For all sequences of states $s_0 s_1 \dots s_n \in S^+$, it must be $\sum_{a \in A(s_n)} Q_\eta(a \mid s_0 s_1 \dots s_n) = 1$. Thus, a policy can be both history-dependent and randomized.

Under policy η , the probability of a transition from s_n to t after the state sequence $s_0 \cdots s_n$ is thus given by $\sum_{a \in A(s_n)} p_{s_n, t}(a) Q_\eta(a \mid s_0 \cdots s_n)$. A policy η gives rise to a probability distribution over the set of behaviors [Der70]. We write $\Pr_s^\eta(\mathcal{A})$ to denote the probability of event \mathcal{A} when policy η is used from the initial state s . We also let X_i and Y_i be the random variables representing the i -th state and the i -th action along a behavior, respectively. We say that a policy η is *memoryless* if $Q_\eta(a \mid s_0 s_1 \cdots s_n) = Q_\eta(a \mid s_n)$ for all sequences of states $s_0 s_1 \cdots s_n \in S^+$ and all $a \in A(s_n)$.

3.2 Probabilistic Fairness

Fairness is a concept that has been introduced in the context of non-probabilistic systems to model the outcome of probabilistic choices while abstracting from the numerical values of the probabilities. Notions of fairness for probabilistic systems have been studied in [HSP83, Var85] and more recently in [KB96], which also present model-checking algorithms for probabilistic systems with fairness.

Given an MDP $\Pi = (S, A, p)$, a *fairness condition* \mathcal{F} for Π is a mapping \mathcal{F} that associates to each $s \in S$ a subset $\mathcal{F}(s) \subseteq A(s)$. The intended meaning is that the choice at s among actions in $\mathcal{F}(s)$ should be “fair.” The various notions of fairness differ in the way in which this “fairness” is defined. According to [KB96], a policy η is said to be *strictly fair* (resp. *almost*, or *weakly*, *fair*) if the behaviors that arise under η all satisfy (resp. satisfy with probability 1) the following condition: *whenever a behavior visits infinitely often a state s , each action in $\mathcal{F}(s)$ is chosen infinitely often at s .* In this paper we introduce a new notion of fairness, called *probabilistic fairness*. Unlike the above notion of fairness, the definition of probabilistic fairness refers directly to the policies, rather than to the behaviors that arise from the policies.

Given an MDP $\Pi = (S, A, p)$ and a fairness condition \mathcal{F} for Π , we say that a policy η is (probabilistically) \mathcal{F} -*fair* if there is $\varepsilon > 0$ such that, for all $n \geq 0$, all sequences of states $s_0, \dots, s_n \in S^+$, and all $a \in \mathcal{F}(s_n)$, we have $Q_\eta(a \mid s_0 \cdots s_n) \geq \varepsilon$. The set of \mathcal{F} -fair policies is denoted by $\eta(\mathcal{F})$.

Clearly, if a policy is \mathcal{F} -fair then it is also weakly fair; the converse is not true in general. In the above definition, ε can depend on the policy η , but cannot depend on the past sequence $s_0 \cdots s_{n-1}$ of states. If ε could depend on the past, then probabilistic fairness would not imply weak fairness. Later we will prove that, for finite TPSs and in the context of the long-run average properties we consider, probabilistic fairness is equivalent to weak fairness. This equivalence does not hold for all types of systems and properties.

A *fair TPS* $\Pi = (S, A, p, S_{in}, time, \mathcal{F})$ consists of a TPS $(S, A, p, S_{in}, time)$ and of a fairness condition \mathcal{F} for (S, A, p) .

3.3 Translating STS into Fair TPS

Given an STS $\mathcal{S} = (\mathcal{V}, \Theta, \mathcal{T})$, its translation TPS $\Pi_{\mathcal{S}} = (S, A, p, S_{in}, time, \mathcal{F})$ shares the same state space S of \mathcal{S} ; the set of initial states is $S_{in} = \{s \in S \mid s \models \Theta\}$. For each $s \in S$, the other components of $\Pi_{\mathcal{S}}$ are defined as follows, depending on whether some immediate transition is enabled at s or not.

3.3.1 Some immediate transition enabled. Let $\mathcal{T}_i(s) = \mathcal{T}(s) \cap \mathcal{T}_i$ be the set of immediate transitions enabled at s , and assume that $\mathcal{T}_i(s) \neq \emptyset$. In this case, we let $A(s) = \{a_\tau \mid \tau \in \mathcal{T}_i(s)\}$, where action a_τ represents the choice of transition τ at s . For all $\tau \in \mathcal{T}_i(s)$, we let $\text{time}(s, a_\tau) = 0$; moreover, action a_τ is fair at s iff τ is fair: precisely, $a_\tau \in \mathcal{F}(s)$ iff $\tau \in \mathcal{T}_f$, for all $\tau \in \mathcal{T}_i(s)$.

For each mode $1 \leq i \leq m_\tau$, action a_τ leads with probability p_i^τ to state $f_i^\tau(s)$, except that if two or more modes lead to the same state, the probabilities are added. Precisely, for all $t \in S$, we let $p_{st}(a_\tau) = \sum_{i=1}^{m_\tau} p_i^\tau \delta[f_i^\tau(s) = t]$, where $\delta[\alpha]$ is 1 if α is true, and 0 otherwise.

3.3.2 No immediate transitions enabled. If $\mathcal{T}(s) \subseteq \mathcal{T}_d$, we let $\mathcal{T}_e(s) = \mathcal{T}(s) \cap \mathcal{T}_e$ and $\mathcal{T}_u(s) = \mathcal{T}(s) \cap \mathcal{T}_u$; note that $\mathcal{T}_e(s) \neq \emptyset$, due to the presence of the idling transition. We let $A(s) = \{a_e\} \cup \{a_\tau \mid \tau \in \mathcal{T}_u(s)\}$: action a_e represents the choice of a transition with exponential distribution, and for $\tau \in \mathcal{T}_u(s)$ action a_τ represents the choice of the transition τ , which has unspecified delay distribution. We let $\mathcal{F}(s) = A(s)$, and we define the expected times of the actions by $\text{time}(s, a_e) = 1 / \sum_{\tau' \in \mathcal{T}_e(s)} \gamma_{\tau'}$, and $\text{time}(s, a_\tau) = 0$ for all $\tau \in \mathcal{T}_u(s)$.

Moreover, for $\tau \in \mathcal{T}_e(s)$ let $p_s(\tau) = \gamma_\tau / \sum_{\tau' \in \mathcal{T}_e(s)} \gamma_{\tau'}$. In other words, $p_s(\tau)$ is the probability that τ is selected at s , conditional to the fact that the transition is selected from $\mathcal{T}_e(s)$. For all $t \in S$ and $\tau \in \mathcal{T}_u(s)$, the transition probabilities are defined by:

$$p_{st}(a_\tau) = \sum_{i=1}^{m_\tau} p_i^\tau \delta[f_i^\tau(s) = t] \quad p_{st}(a_e) = \sum_{\tau \in \mathcal{T}_e(s)} \sum_{i=1}^{m_\tau} p_\tau(s) p_i^\tau \delta[f_i^\tau(s) = t].$$

3.4 Non-Zeno TPSs

We say that a fair TPS is *non-Zeno* if time diverges with probability 1 along all behaviors, under all fair policies. Precisely, $\Pi = (S, A, p, S_{in}, \text{time}, \mathcal{F})$ is non-Zeno iff we have $\Pr_s^\eta(\sum_{k=0}^\infty \text{time}(X_k, Y_k) = \infty) = 1$ for all $s \in S_{in}$ and all $\eta \in \boldsymbol{\eta}(\mathcal{F})$. Since behaviors along which time does not diverge have no physical meaning, we only consider non-Zeno TPSs: after translating an STS into a fair TPS, it is necessary to check that it is non-Zeno. A method to do this is presented in Section 6. A more sophisticated approach to the problem of time divergence, inspired by [Seg95], is discussed in [dA97a].

4 Translation and Informal Semantics

Even though the formal semantics of STSs is defined by translation into fair TPSs, there is a correspondence between the proposed translation and the informal semantics presented in Section 2.1. This correspondence is important from a pragmatic point of view, since system models are usually constructed with this intuitive semantics in mind. We justify the translation in three steps, considering first the structure of the translation TPS, then the use of fairness, and lastly the interaction between translation and specification languages.

4.1 Structure of the Translation TPS

To understand the correspondence between the translation and the informal semantics, consider the system evolution from a state s . If there are immediate transitions enabled at s , the correspondence between the informal semantics and the structure of the translation TPS is immediate.

If $\mathcal{T}(s) \subseteq \mathcal{T}_d$, let as before $\mathcal{T}_e(s) = \mathcal{T}(s) \cap \mathcal{T}_e$ and $\mathcal{T}_u(s) = \mathcal{T}(s) \cap \mathcal{T}_u$. The set of available actions at s is $\{a_e\} \cup \{a_\tau \mid \tau \in \mathcal{T}_u(s)\}$. Let q_e and q_τ , for $\tau \in \mathcal{T}_u(s)$, be the probabilities with which these actions are chosen by a policy. Note that q_e and q_τ can depend on the past history of the behavior. There is a relation between the probabilities q_e and q_τ , $\tau \in \mathcal{T}_u$, selected by the policy, and the rates of the transitions in $\mathcal{T}_u(s)$, selected nondeterministically in the informal semantics. To derive the relation, consider the probability of choosing $\tau \in \mathcal{T}(s)$ in the translation TPS and in the informal semantics. In the TPS, this probability is equal to q_τ for $\tau \in \mathcal{T}_u(s)$, and to $q_e p_\tau(s)$ for $\tau \in \mathcal{T}_e(s)$. In the informal semantics this probability is equal to $\gamma_\tau / \sum_{\tau' \in \mathcal{T}(s)} \gamma_{\tau'}$ for all $\tau \in \mathcal{T}(s)$. Equating these probabilities, we obtain

$$q_e = \left(\sum_{\tau' \in \mathcal{T}_e(s)} \gamma_{\tau'} \right) / \left(\sum_{\tau' \in \mathcal{T}(s)} \gamma_{\tau'} \right) \quad q_\tau = \gamma_\tau / \left(\sum_{\tau' \in \mathcal{T}(s)} \gamma_{\tau'} \right) \quad (1)$$

for all $\tau \in \mathcal{T}_u(s)$. This relation between q_e , $\{q_\tau\}_{\tau \in \mathcal{T}_u(s)}$ and $\{\gamma_\tau\}_{\tau \in \mathcal{T}_u(s)}$ preserves not only the probabilities of the transitions from s , but also the expected time spent at s . In fact, from Section 3.3.2 the expected time spent by the TPS at s is equal to $q_e / \sum_{\tau' \in \mathcal{T}_e(s)} \gamma_{\tau'}$. If we substitute into this equation the value of q_e given by (1), we obtain $1 / \sum_{\tau' \in \mathcal{T}(s)} \gamma_{\tau'}$, which is exactly the expected time spent at s under the informal semantics. Thus, equations (1) together with the constraint $q_e + \sum_{\tau \in \mathcal{T}_u(s)} q_\tau = 1$ define a one-to-one mapping between the unspecified transition rates in the informal semantics and the probabilities of choosing the actions in the translation TPS. The mapping preserves both the expected time spent at s , and the probabilities of transitions from s . Given a nondeterministic choice for the transition rates $\{\gamma_\tau\}_{\tau \in \mathcal{T}_u(s)}$, we can determine a policy which simulates this choice; conversely, each policy can be interpreted as a choice for these rates. This correspondence indicates that the translation from STSs to fair TPS preserves the informal semantics of STSs.

4.2 Translation and Fairness

The above considerations also justify our use of fairness in the translation. In fact, for $\tau \in \mathcal{T}_u(s)$ the fairness of a_τ requires that $q_\tau > 0$, which by (1) corresponds to the requirement $\gamma_\tau > 0$. Similarly, the fairness of a_e requires that $q_e > 0$, which corresponds to the requirement $\gamma_\tau < \infty$ for all $\tau \in \mathcal{T}_u(s)$. Thus, the fairness conditions and the notion of probabilistic fairness are the exact counterpart of the requirements $0 < \gamma_\tau < \infty$ for the rates of transitions $\tau \in \mathcal{T}_u$.

4.3 Translation and Specification Language

In Section 3.3.2 we assign expected time 0 to the actions that correspond to transitions with unspecified rates. The argument presented above to justify this assignment is valid only if we assume a restriction on the expressive power of specification methods. Precisely, we allow specification methods to refer to the amount of time spent at a state, but we require that they do not measure this amount of time *conditional on the successor state*.

To clarify this point, consider as an example a state s of an STS on which two transitions are enabled: a transition τ_1 with rate γ_1 , leading to state t_1 , and a transition τ_2 with unspecified rate, leading to state t_2 . The translation we proposed would be inappropriate if our specification methods could express properties like: “*the time spent at s when t_1 is the immediate successor is on average $> b$.*” In fact, for the purposes of this property the choice of assigning $time(s, a_{\tau_2}) = 0$ would not correspond to the idea of assigning nondeterministically a transition rate to τ_2 . On the other hand, if the specification methods can refer only to the expected time spent at s , regardless of the successor of s , then the translation is faithful to the informal semantics. The specification methods discussed in the next section obey this restriction.

5 Specification of Long-Run Average Properties

The long-run average properties we consider in this paper refer to the average outcome of a *task*, measured over an interval of time whose length diverges to infinity. A *task* is a (hopefully) finite activity performed regularly by the system. The outcome of the task can depend both on its completion, and on its duration. For example, a task might consist in sending a message and waiting for the acknowledge; its outcome might be 1 if the acknowledge is received, or 0 if a timeout occurs. The long-run average outcome of this task is equal to the long-run average fraction of messages that are acknowledged. In [dA97a], tasks were specified using labeled graphs called *experiments*. Here, we follow a simplified approach, and given a fair TPS $\Pi = (S, A, p, S_{in}, time, \mathcal{F})$, we specify tasks and their outcomes using two labelings w and r :

- The labeling $w : S \times S \mapsto \{0, 1\}$ associates to each $s, t \in S$ a label $w(s, t)$, which has value 1 if the transition $s \rightarrow t$ completes a task, and 0 otherwise.
- The labeling $r : S \times \bigcup_{s \in S} A(s) \times S \mapsto \mathbb{R}^+$ is used to define the outcome of a task. Due to the restrictions on specification languages mentioned in Section 4.3, we consider only labelings that can be written in the form

$$r(s, a, t) = \alpha(s) time(s, a) + \beta(s, t)$$

for some functions $\alpha : S \mapsto \mathbb{R}^+$ and $\beta : S \times S \mapsto \mathbb{R}^+$ (where $\mathbb{R}^+ = \{x \in \mathbb{R} \mid x \geq 0\}$). Thus, the labeling r can be used to measure the expected time spent at system states, weighted by a function α ; the “cost” associated to system transitions, expressed by β ; or a combination of the two.

In GSPN reward models [CMT91] it is possible to associate a reward rate to the places and transitions of the net; [Cla96] and [Ber97] propose methods for associating a reward with each state of the Markov chain generated from a PEPA or EMPA model. The r labeling discussed above serves a similar purpose; however, we also introduce the notion of *task*, and the corresponding w labeling. For systems that can be translated into ergodic Markov chains, the two approaches are equally expressive: even without a w labeling, the average outcome of a task can be measured by measuring separately the rates of task completion and of output generation, and by taking the ratio between the two. In the case of systems with nondeterministic behavior, however, our approach leads to more expressive specification methods. In fact, in these systems the choice of policy may influence differently the task completion rate and the outcome generation rate. Thus, the ratio between the maximal outcome generation rate and the minimal task completion rate is in general *not* equal to the maximal long-run average outcome of a task. From the r , w labelings, for each behavior ω of Π we define a predicate I and a quantity H_n , for $n \geq 0$, as follows:

$$I \text{ iff } \exists^\infty k . [r(X_k, Y_k, X_{k+1}) > 0 \vee w(X_k, X_{k+1}) > 0] \quad (2)$$

$$H_n = \frac{\sum_{k=0}^{n-1} r(X_k, Y_k, X_{k+1})}{\sum_{k=0}^{n-1} w(X_k, X_{k+1})} . \quad (3)$$

In (2), the notation $\exists^\infty k$ is an abbreviation for “there are infinitely many distinct values for k ”. Thus, I holds if ω completes infinitely many tasks, or if one such task produces infinite outcome. The quantity H_n represents the average outcome per task for the first n steps of ω . For all $s \in S$ and all policies η , we let

$$H_\eta^-(s) = \inf \{a \in \mathbb{R} \mid \Pr_s^\eta(I \wedge \liminf_{n \rightarrow \infty} H_n \leq a) > 0\} \quad (4)$$

be the infimum of the set of long-run average outcomes obtained with non-zero probability by behaviors that satisfy I . We do not consider behaviors on which I is false, since these behaviors after a certain position cease to complete tasks or to produce outcome, and the long-run average outcome is consequently not well-defined: this point is discussed in detail in [dA97a, dA98]. Finally, we let

$$H_{\mathcal{F}}^-(s) = \inf_{\eta \in \boldsymbol{\eta}(\mathcal{F})} H_\eta^-(s) \quad H_{\mathcal{F}}^+(s) = \sup_{\eta \in \boldsymbol{\eta}(\mathcal{F})} H_\eta^+(s) .$$

The quantities $H_{\mathcal{F}}^-(s)$ and $H_{\mathcal{F}}^+(s)$ represent the minimal and maximal long-run average outcomes that can be achieved with non-zero probability by I -behaviors, provided that the long-run average outcome is well-defined, and that a \mathcal{F} -fair policy is used from s . The specification of long-run average properties of STSs and fair STSs is based on the specification of lower (resp. upper) bounds for $H_{\mathcal{F}}^-(s)$ (resp. $H_{\mathcal{F}}^+(s)$), for some states $s \in S$.

As an example, consider the commuter system of Section 2.2. For all $s, t \in S$, we let $w(s, t) = 1$ if $s \llbracket c \rrbracket = 4$ and $t \llbracket c \rrbracket = 1$, and $w(s, t) = 0$ otherwise, so that w counts the number of returns to city A. For all $s \in S$ and $a \in A(s)$, we also let

$r(s, a) = \text{time}(s, a)$ if $s \llbracket c \rrbracket \in \{1, 2\}$, and $r(s, a) = 0$ otherwise, so that r measures the time spent going from A to B. With these labelings, $H_{\mathcal{F}}^+(s)$ is equal to the maximal long-run average duration of a one-way trip from city A to city B, if the system is initially at s (it can be shown that $H_{\mathcal{F}}^+(s)$ does not depend on s in this system). Using the algorithm presented in Section 6, we can compute that $H_{\mathcal{F}}^+(s) \simeq 7.5526$.

6 Verification of Long-Run Average Properties

The verification problem for long-run average properties consists in computing $H_{\mathcal{F}}^+(s)$, $H_{\mathcal{F}}^-(s)$ at all states $s \in S$ of a fair TPS. Algorithms that solve this verification problem for the case without fairness conditions have been presented in [dA97a, dA98]. To solve the model-checking problem in presence of fairness conditions, we first decompose the fair TPS into the components where a behavior can reside forever under a fair policy. These components are called *fair end components*, and are presented below. Once the TPS has been decomposed, we apply to each component the algorithm of [dA98] to compute the maximal and minimal long-run average outcome for the component, disregarding the fairness conditions. These maximal and minimal values correspond to optimal and pessimal policies, which need not be fair. Nevertheless, using results on *parametric Markov chains* we show that we can approximate these policies with a series of fair policies, whose long-run average outcome converges to that of the optimal and pessimal policies. This shows that, for each component, the maximal and minimal long-run average outcomes computed disregarding fairness conditions also apply to the case with fairness. Hence, the values of $H_{\mathcal{F}}^+(s)$ and $H_{\mathcal{F}}^-(s)$ at a state s can be obtained by taking the maximum and minimum values of the long-run average outcome computed for any component reachable from s .

6.1 Fair End Components

Given an MDP $\Pi = (S, A, p)$, a *sub-MDP* is a pair (C, D) , where $C \subseteq S$ and D is a function that associates to each $s \in C$ a subset $D(s) \subseteq A(s)$ of actions. The sub-MDP corresponds thus to a subset of states and actions of the original MDP. We say that a sub-MDP (C, D) is contained in a sub-MDP (C', D') if $\{(s, a) \mid s \in C \wedge a \in D(s)\} \subseteq \{(s, a) \mid s \in C' \wedge a \in D'(s)\}$.

Given a fairness condition \mathcal{F} for Π , we say that sub-MDP (C, D) is a *fair end component* (FEC) if the following conditions hold [dA97a]:

- *Closure*: for all $s \in C$, $a \in D(s)$, and $t \in S$, if $p_{st}(a) > 0$ then $t \in C$.
- *Connectivity*: Let $E = \{(s, t) \in C \times C \mid \exists a \in D(s) . p_{st}(a) > 0\}$. The graph (C, E) is strongly connected.
- *Fairness*: For all $s \in C$, we have $\mathcal{F}(s) \subseteq D(s)$.

We say that a FEC (C, D) is *maximal* if there is no other FEC (C', D') that properly contains (C, D) . We denote by $\text{MFEC}(\Pi, \mathcal{F})$ the set of maximal FECs of Π . The set $\text{MFEC}(\Pi, \mathcal{F})$ can be computed in time polynomial in $\sum_{s \in S} |A(s)|$ using simple graph algorithms; an algorithm to do so is given in [dA97a, §8].

Intuitively, a fair end component is a portion of MDP consisting of the states and actions that can be visited infinitely often by a behavior with positive probability, under some fair policy. To make this concept precise, given a behavior ω we let $(C, D) = \text{inft}(\omega)$ be the sub-MDP defined by $C = \{s \mid \exists^\infty k. X_k = s\}$ and, for $s \in C$, $D(s) = \{a \mid \exists^\infty k. X_k = s \wedge Y_k = a\}$.

Theorem 1 *For any $s \in S$ and $\eta \in \eta_{\mathcal{F}}$, we have $\text{Pr}_s^\eta(\text{inft}(\omega) \text{ is a FEC}) = 1$.*

In a purely probabilistic system, fair end components correspond to the closed recurrent classes of the Markov chain underlying the system [KSK66]. Fair end components are the fair counterpart of the end components of [dA97a, dA98], and are related to sets used in [KB96] to solve the model-checking problem for PCTL*. As our first application of the above theorem, we obtain a criterion to decide whether a fair TPS is non-Zeno.

Theorem 2 (condition for non-Zenoness) *Given a fair TPS $\Pi = (S, A, p, S_{in}, \text{time}, \mathcal{F})$, a FEC (C, D) is a zero-FEC if $\text{time}(s, a) = 0$ for all $s \in C$ and $a \in D(s)$. TPS Π is non-Zeno iff there is no zero-FEC reachable from S_{in} .*

Even though there can be exponentially many zero-FECs in a fair TPS, it is easy to see that it suffices to consider the maximal ones. Hence, checking non-Zenoness can be done in time polynomial in $\sum_{s \in S} |A(s)|$ [dA97a, §8].

6.2 Parametric Markov Chains

Given a finite set S of indices, a *substochastic matrix* is a matrix $P = [p_{st}]_{s,t \in S}$ such that $0 \leq p_{st} \leq 1$ for all $s, t \in S$ and $\sum_{t \in S} p_{st} \leq 1$ for all $s \in S$. Given a sub-stochastic matrix P , the *steady-state distribution matrix* is defined by $P^* = \lim_{n \rightarrow \infty} n^{-1} \sum_{k=0}^{n-1} P^k$ [KSK66]. We say that a state of P is *surely recurrent* if the Markov chain corresponding to P has only one closed recurrent class, and if the state belongs to that class. The following result can be proved by linear algebra arguments [dA97a, §8].

Theorem 3 (continuity of steady-state distributions) *Consider a family $P(x) = [p_{st}(x)]_{s,t \in S}$ of substochastic matrices parameterized by $x \in I$, where $I \subseteq \mathbb{R}$ is an interval of real numbers. Assume that the coefficients of $P(x)$ depend continuously on x for $x \in I$. If there is $s \in S$ that is surely recurrent for all $x \in I$, then also the coefficients of $P^*(x)$ depend continuously on x for $x \in I$.*

6.3 The Model-Checking Algorithm

From the definitions of $H_{\mathcal{F}}^-(s)$ and $H_{\mathcal{F}}^+(s)$, we see that these quantities depend only on the states and actions that are repeated infinitely often. Theorem 1 states that these states and actions form a FEC with probability 1: hence, we can concentrate our attention on the maximal FECs. We say that an MDP is *strongly connected* if, for each pair of states, there is a behavior prefix that leads from one state to the other. By definition, maximal FECs are strongly connected

sub-MDPs. Denote by $\emptyset = \lambda s. \emptyset$ the empty fairness condition. The following theorem summarizes several results of [dA97a, §5] for strongly connected MDPs without fairness conditions.

Theorem 4 *Consider a strongly connected TPS $\Pi = (S, A, p, r, w)$. The following assertions hold.*

- *The value of $H_{\emptyset}^{-}(s)$ does not depend on $s \in S$. The common value H_{\emptyset}^{-} can be computed in time polynomial in the size of Π .*
- *There is a memoryless policy η such that $H_{\eta}^{-}(s) = H_{\emptyset}^{-}$ for all $s \in S$. Moreover, the transition matrix $P_{\eta} = [p_{st}^{\eta}]_{s,t \in S}$ defined by $p_{st}^{\eta} = \sum_{a \in A(s)} p_{st}(a) Q_{\eta}(a \mid s)$ corresponds to a Markov chain having a single closed recurrent class.*

Similar assertions hold for $H_{\emptyset}^{+}(s)$.

Using the results of the above theorem, we propose the following algorithms for the computation of $H_{\mathcal{F}}^{-}(s)$ and $H_{\mathcal{F}}^{+}(s)$.

Algorithm 1 (computation of $H_{\mathcal{F}}^{-}(s)$ and $H_{\mathcal{F}}^{+}(s)$) Given a fair TPS $\Pi = (S, A, p, S_{in}, time, \mathcal{F})$ together with labelings r, w , the quantities $H_{\mathcal{F}}^{-}(s)$ and $H_{\mathcal{F}}^{+}(s)$ can be computed at all $s \in S$ as follows.

1. Let $\mathcal{L} = \{(C, D) \in \text{MFEC}(\Pi, \mathcal{F}) \mid \exists s, t \in C. \exists a \in D(s). [r(s, a, t) > 0 \vee w(s, t) > 0]\}$ be the set of maximal FECs that contain at least one instance of strictly positive r or w label. Write $\mathcal{L} = \{(C_1, D_1), \dots, (C_n, D_n)\}$.
2. For each $1 \leq i \leq n$, let $\Pi_i = (C_i, D_i, p^i, \mathcal{F}_i, r_i, w_i)$, where $p^i, \mathcal{F}_i, r_i, w_i$ are the restrictions of p, \mathcal{F}, r, w to C_i, D_i , for $1 \leq i \leq n$. Using Theorem 4, compute the values $H_{\emptyset, i}^{-}, H_{\emptyset, i}^{+}$ for all MDPs Π_i , for $1 \leq i \leq n$.
3. For each $s \in S$, let $K(s) = \{i \in [1..n] \mid s \text{ can reach } C_i\}$ be the set of indices of maximal FECs reachable from s . Then, $H_{\mathcal{F}}^{-}(s) = \min_{i \in K(s)} H_{\emptyset, i}^{-}$ and $H_{\mathcal{F}}^{+}(s) = \min_{i \in K(s)} H_{\emptyset, i}^{+}$. ■

Theorem 5 *Algorithm 1 correctly computes $H_{\mathcal{F}}^{-}(s)$ and $H_{\mathcal{F}}^{+}(s)$.*

Proof (sketch). The crux of the argument is to show that in a strongly connected MDP the equality $H_{\mathcal{F}}^{-}(s) = H_{\emptyset}^{-}$ holds for all s (and similarly for $H_{\mathcal{F}}^{+}(s)$). Once this is done, the decomposition in maximal FECs (Step 1) is justified by Theorem 1, and the selection of the maximal FECs that contain at least one positive r, w label is justified by (2), (3) and (4). Finally, Step 3 can be justified using simple reachability arguments.

To show that in a strongly connected MDP $\Pi = (S, A, p, \mathcal{F}, r, w)$ we have $H_{\mathcal{F}}^{-}(s) = H_{\emptyset}^{-}$ for all s , it suffices to show that $H_{\mathcal{F}}^{-}(s) = H_{\eta^*}^{-}(s)$, where η^* is the policy described in Theorem 4. To this end, let η^{\bullet} be the memoryless \mathcal{F} -fair policy that at each $s \in S$ chooses uniformly at random an action from $A(s)$. For each $0 \leq x < 1$, define the memoryless policy $\eta(x)$ by

$$Q_{\eta(x)}(a \mid s) = x Q_{\eta^{\bullet}}(a \mid s) + (1 - x) Q_{\eta^*}(a \mid s)$$

for all s and all $a \in A(s)$. Note that policy $\eta(x)$ is \mathcal{F} -fair for $0 < x < 1$, and it coincides with η^* for $x = 0$. Let $P(x) = [p_{st}(x)]_{s,t \in S}$ be the matrix of the Markov chain arising from $\eta(x)$, defined by

$$p_{st}(x) = \sum_{a \in A(s)} Q_{\eta(x)}(a | s) p_{st}(a) ,$$

and let

$$r_s(x) = \sum_{a \in A(s)} \sum_{t \in S} Q_{\eta(x)}(a | s) p_{st}(a) r(s, a, t) \quad w_s(x) = \sum_{t \in S} p_{st}(x) w(s, t) ,$$

for all $s \in S$ and $0 \leq x < 1$. Denote by $P^*(x) = [p_{st}^*]_{s,t \in S}$ the steady-state distribution matrix corresponding to $P(x)$. By our choice of η^* (see Theorem 4), the Markov chain corresponding to $P(0)$ has a single closed recurrent class $C \subseteq S$. Since the MDP is strongly connected, by definition of $\eta(x)$ all states of C are surely recurrent for $0 \leq x < 1$. Hence, as a consequence of standard facts on Markov chains we have $H_{\eta(x)}^-(s) = (\sum_{t \in S} p_{st}^* r_t) / (\sum_{t \in S} p_{st}^* w_t)$. Theorem 3 ensures that $\lim_{x \rightarrow 0} P^*(x) = P^*(0)$. Since for all $s \in S$ quantities $r_s(x)$ and $w_s(x)$ are also continuous for $x \rightarrow 0$, we have $\lim_{x \rightarrow 0} H_{\eta(x)}^-(s) = H_{\emptyset}^-$. From $H_{\emptyset}^- \leq H_{\mathcal{F}}^-(s)$ and from the fact that $\eta(x)$ is \mathcal{F} -fair follows $H_{\mathcal{F}}^-(s) = \inf_{\eta \in \boldsymbol{\eta}(\mathcal{F})} H_{\eta}^-(s) = H_{\emptyset}^-$, as was to be proved. ■

The complexity of the model-checking problem is given by the following result, which is an immediate consequence of Theorem 4 and Algorithm 1.

Theorem 6 *The complexity of the model-checking problem for $H_{\mathcal{F}}^-(s)$, $H_{\mathcal{F}}^+(s)$ is polynomial in the size of the translation TPS.*

We conclude by showing that the notions of weak fairness and probabilistic fairness coincide for finite TPSs and long-run average properties.

Theorem 7 *Let $\tilde{H}_{\mathcal{F}}^-(s) = \inf_{\eta \in \tilde{\boldsymbol{\eta}}(\mathcal{F})} H_{\eta}^-(s)$, where $\tilde{\boldsymbol{\eta}}(\mathcal{F})$ is the set of weakly fair policies, defined according to [KB96]. Then, $H_{\mathcal{F}}^-(s) = \tilde{H}_{\mathcal{F}}^-(s)$. A similar result holds for $H_{\mathcal{F}}^+(s)$.*

Proof. The result follows from an analysis of the proof of Theorem 5, together with the observation that probabilistically fair policies are also weakly fair. ■

Acknowledgments. We thank Rob van Glabbeek for helpful discussions on the translation from STSs to TPSs, and the anonymous referees for useful comments.

References

- [ABC84] M. Ajmone Marsan, G. Balbo, and G. Conte. A class of generalized stochastic Petri nets for the performance analysis of multiprocessor systems. *ACM Trans. Comp. Sys.*, 2(2):93–122, 1984.

- [ACD92] R. Alur, C. Courcoubetis, and D. Dill. Verifying automata specifications of probabilistic real-time systems. In *Real Time: Theory in Practice*, volume 600 of *LNCS*, pages 28–44. Springer-Verlag, 1992.
- [Ber97] M. Bernardo. An algebra-based method to associate rewards with EMPA terms. In *Proc. ICALP'97*, volume 1256 of *LNCS*, pages 358–368. Springer-Verlag, 1997.
- [BG96] M. Bernardo and R. Gorrieri. Extended Markovian process algebra. In *Proc. CONCUR'96*, volume 1119 of *LNCS*, pages 315–330. Springer-Verlag, 1996.
- [CMT91] G. Ciardo, J.K. Muppala, and K.S. Trivedi. On the solution of GSPN reward models. *Performance Evaluation*, 12:237–253, 1991.
- [Cla96] G. Clark. Formalising the specification of rewards with PEPA. In *Proc. 4th Workshop on Process Algebras and Performance Modelling*, pages 139–160. CLUT, Torino, Italy, 1996.
- [dA97a] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997. Technical Report STAN-CS-TR-98-1601.
- [dA97b] L. de Alfaro. Temporal logics for the specification of performance and reliability. In *Proc. STACS'97*, volume 1200 of *LNCS*, pages 165–176. Springer-Verlag, 1997.
- [dA98] L. de Alfaro. How to specify and verify the long-run average behavior of probabilistic systems. In *Proc. LICS'98*, 1998. To appear.
- [Der70] C. Derman. *Finite State Markovian Decision Processes*. Academic Press, 1970.
- [GHR93] H.N. Götz, U. Herzog, and M. Rettelsbach. Multiprocessor and distributed system design: the integration of functional specification and performance analysis using stochastic process algebras. In *PERFORMANCE'93*, volume 729 of *LNCS*, Springer-Verlag, 1993.
- [Hil96] J. Hillston. *A Compositional Approach to Performance Modelling*. Distinguished Dissertations Series. Cambridge University Press, 1996.
- [HSP83] S. Hart, M. Sharir, and A. Pnueli. Termination of probabilistic concurrent programs. *ACM Trans. Prog. Lang. Sys.*, 5(3):356–380, 1983.
- [KB96] M. Kwiatkowska and C. Baier. Model checking for a probabilistic branching time logic with fairness. To appear in *Distributed Computing*. Preliminary version in Technical Report CSR-96-12, University of Birmingham, 1996.
- [KSK66] J.G. Kemeny, J.L. Snell, and A.W. Knapp. *Denumerable Markov Chains*. D. Van Nostrand Company, 1966.
- [MP91] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, New York, 1991.
- [PZ86] A. Pnueli and L. Zuck. Probabilistic verification by tableaux. In *Proc. LICS'86*, pages 322–331, 1986.
- [Seg95] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995. Technical Report MIT/LCS/TR-676.
- [SL94] R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. In *Proc. CONCUR'94*, volume 836 of *LNCS*, pages 481–496. Springer-Verlag, 1994.
- [Var85] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *Proc. FOCS'85*, pages 327–338, 1985.